Technical Note: Context Bandwidth in Context-Driven Learning Systems

Core Idea

Context bandwidth is the amount of contextual information available to select or modulate computation at a particular stage of a model.

Rather than asking:

- How many parameters does the model have?
- How deep is the model?

we can also ask:

- How much context is available?
- How efficiently is that context exploited?

Context bandwidth provides a common language for comparing context trees, ReLU neural networks, and Atlas neural networks, despite their very different implementations.

Three Examples

Context Trees

A context tree consumes context **incrementally**.

Each decision contributes a small amount of information before selecting the next branch.

Characteristics:

- Narrow context bandwidth.
- Sequential context accumulation.
- Local predictors reached after relatively few decisions.
- Limited ability to model high-order interactions without very deep trees.

The amount of context available to the predictor is fundamentally limited by the depth of the tree.

ReLU Neural Networks

A ReLU network constructs context from the activation state of many neurons simultaneously.

Each layer generates a large binary activation pattern that determines which local linear computation is active.

Characteristics:

- Very high context bandwidth.
- Massive parallel context representation.
- Rich interactions between context variables.
- Computation factorized through shared linear mappings.

Although many activation bits are correlated, the available context can still be extremely information-rich compared with classical context trees.

Atlas Neural Networks

Atlas neural networks explicitly separate context generation from prediction.

The context may be generated from methods such as:

- random projections,
- LSH,
- atan2 wedge partitioning,
- learned binary propositions,
- other geometric indexing schemes.

This explicit context is then used to select parameter subsets, routing paths, or local predictors.

Characteristics:

- Explicit high-bandwidth context.
- Context representation is directly observable.
- Independent design of context generation and local predictors.
- Enables statistical analysis of the context itself.

Unlike conventional neural networks, the context representation becomes a first-class object that can be measured, analysed, and optimized independently of the predictor.

Effective Context Bandwidth

Raw context dimensionality is not the same as effective bandwidth.

The useful bandwidth depends on properties such as:

- entropy,
- redundancy,
- mutual information with the prediction target,
- stability under perturbations,
- frequency of context occurrence.

Many nominal context bits may contribute little useful information if they are highly correlated.

Bandwidth Utilization

Context bandwidth alone does not determine model quality.

Important questions include:

- How much of the available context actually influences prediction?
- Which context variables are most informative?
- How effectively are parameters shared across contexts?
- How much computation is reused?

Efficient systems maximize the predictive value obtained from each context bit.

Context Geometry

Bandwidth should be considered together with the geometry of the context space.

Useful questions include:

- Are nearby contexts semantically related?
- Do small input perturbations produce small context changes?
- How uniformly is context space occupied?
- How fragmented are occupied regions?
- Does the context preserve neighbourhood structure?

The geometry determines how smoothly computation changes across neighbouring contexts.

Architectural Perspective

The three architectures illustrate different design philosophies.

Context Trees

- Small context bandwidth.
- Strong specialization.
- Explicit logical partitioning.

ReLU Neural Networks

- Large implicit context bandwidth.
- Extensive parameter sharing.
- Geometric factorization through shared linear mappings.

Atlas Neural Networks

- Large explicit context bandwidth.
- Context generation separated from prediction.
- Context treated as an observable computational state.

- Statistical characterization of context becomes possible during both training and inference.

Research Questions

Viewing learning systems through the lens of context bandwidth suggests several research directions:

- What context bandwidth is required for a given task?
- How much effective bandwidth survives after removing redundancy?
- Can bandwidth be allocated adaptively across layers?
- How should predictor complexity scale with context bandwidth?
- Can context quality be measured independently of predictor quality?
- What trade-off exists between explicit and implicit context representations?

Summary

Context bandwidth provides a complementary architectural metric alongside depth, width, and parameter count. It measures the information available to control computation rather than the computation itself.

Viewed from this perspective:

- Context trees operate with narrow, sequentially accumulated contexts.
- ReLU neural networks operate with large implicit contexts that are exploited through geometric factorization and parameter sharing.
- Atlas neural networks make these high-bandwidth contexts explicit, allowing them to be analysed, optimized, and controlled independently of the local predictors.

This viewpoint suggests that context bandwidth, effective bandwidth, and bandwidth utilization may become useful quantitative tools for comparing and designing future context-driven learning architectures.